

## **BACKGROUND OF THE INVENTION**

### **[001] Field of the Invention**

5      **[002]** This invention relates to computer system reliability and, more particularly, to the detection of errors in memory subsystems.

### **[003] Description of the Related Art**

10     **[004]** Computer systems are typically available in a range of configurations which may afford a user varying degrees of reliability, availability and serviceability (RAS). In some systems, reliability may be paramount. Thus, a reliable system may include features designed to prevent failures. In other systems, availability may be important and so systems may be designed to have significant fail-over capabilities in the event of a failure.

15     Either of these types of systems may include built-in redundancies of critical components. In addition, systems may be designed with serviceability in mind. Such systems may allow fast system recovery during system failures due to component accessibility. In critical systems, such as high-end servers and some multiple processor and distributed processing systems, a combination of the above features may produce the  
20     desired RAS level.

25     **[005]** Depending on the type of system, data that is stored in system memory may be protected from corruption in one or more ways. One such way to protect data is to use error detection and/ or error correction codes (ECC). The data may be transferred to system memory with an associated ECC code which may have been generated by a sending device. ECC logic may then regenerate and compare the ECC codes prior to storing the data in system memory. When the data is read out of memory, the ECC codes

may again be regenerated and compared with the existing codes to ensure that no errors have been introduced to the stored data.

[006] In addition, some systems may employ ECC codes to protect data that is routed through out the system. However, in systems where a system memory module such as for example, a dual in-line memory module (DIMM) is coupled to a memory controller, the data bus and corresponding data may be protected as described above but the address, command and control information and corresponding wires may not. In such systems, a bad bit or wire which conveys erroneous address or command information may be undetectable as such an error. For example, correct data may be stored to an incorrect address or data may not be actually written to a given location. When the data is read out of memory, the ECC codes for that data may not detect this type of error, since the data itself may be good. When a processor tries to use the data however, the results may be unpredictable or catastrophic.

15

## SUMMARY OF THE INVENTION

[007] Various embodiments of a memory subsystem are disclosed. In one embodiment, 5 a memory subsystem includes a memory controller coupled to a memory module including a plurality of memory chips via a memory bus. The memory controller may generate a plurality of memory requests each including address information and corresponding error detection information. The corresponding error detection information may be dependent upon the address information. A memory module may 10 receive each of the plurality of memory requests. An error detection circuit within the memory module may detect an error in the address information based upon the corresponding error detection information and may provide an error indication in response to detecting the error.

15 [008] In another embodiment, a memory subsystem includes a memory controller coupled to a memory module including a plurality of memory chips via a memory bus. The memory controller may generate a plurality of memory requests each including control information and corresponding error detection information. The corresponding error detection information may be dependent upon the control information. A memory 20 module may receive each of the plurality of memory requests. An error detection circuit within the memory module may detect an error in the control information based upon the corresponding error detection information and may provide an error indication in response to detecting the error.

**BRIEF DESCRIPTION OF THE DRAWINGS**

[009] FIG. 1 is a block diagram of one embodiment of a computer system.

5

[010] FIG. 2 is a block diagram of one embodiment of a memory subsystem.

[011] FIG. 3 is a block diagram of one embodiment of a memory module.

10 [012] While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and  
15 alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

## DETAILED DESCRIPTION OF THE EMBODIMENTS

[013] Turning now to FIG. 1, a block diagram of one embodiment of a computer system 10 is shown. Computer system 10 includes a plurality of processors 20-20n connected to a memory subsystem 50 via a system bus 25. Memory subsystem 50 includes a memory controller 30 coupled to a system memory 40 via a memory bus 35. It is noted that, although two processors and one memory subsystem are shown in FIG. 1, embodiments of computer system 10 employing any number of processors and memory subsystems are contemplated. In addition, elements referred to herein with a particular reference number followed by a letter may be collectively referred to by the reference number alone. For example, processor 20A-n may be collectively referred to as processor 20.

[014] Memory subsystem 30 is configured to store data and instruction code within system memory 40 for use by processor 20. As will be described further below, in one embodiment, system memory 40 may be implemented using a plurality of dual in-line memory modules (DIMM). Each DIMM may employ a plurality of random access memory chips such as dynamic random access memory (DRAM) or synchronous dynamic random access memory (SDRAM) chips, for example. Although it is contemplated that other types of memory may be used. Each DIMM may be mated to a system memory board via an edge connector and socket arrangement. The socket may be located on a memory subsystem circuit board and each DIMM may have an edge connector which may be inserted into the socket, for example.

[015] Generally speaking, processor 20 may access memory subsystem 50 by initiating a memory request transaction such as a memory read or a memory write to memory controller 30 via system bus 25. Memory controller 30 may then control the storing to and retrieval of data from system memory 40 by issuing memory request commands to

system memory 40 via memory bus 35. Memory bus 35 conveys address and control information and data to system memory 40. The address and control information may be conveyed to each DIMM in a point-to-multipoint arrangement while the data may be conveyed directly between each memory chip on each DIMM in a point-to-point

5 arrangement. The point-to-multipoint arrangement is sometimes referred to as a multi-drop topology.

[016] Referring to FIG. 2, a block diagram of one embodiment of a memory subsystem is shown. Circuit components that correspond to components shown in FIG. 1 are numbered identically for clarity and simplicity. In FIG. 2, memory subsystem 50 includes a memory controller 30 coupled to a system memory 40 via a memory bus 35. Memory controller 30 includes a memory control logic unit 31 and an error detection generation circuit 32. In addition to memory bus 35, two additional signals are conveyed between memory controller 30 and system memory 40: error detection information 36 and error indication 37. As mentioned above, system memory 40 includes a plurality of memory modules depicted as memory modules 0 through n, where n is representative of any number of memory modules.

[017] In the illustrated embodiment, memory controller 30 may receive a memory request via system bus 25. Memory controller logic 31 may then schedule the request and generate a corresponding memory request for transmission on memory bus 35. The request may include address and control information. For example, if the memory request is a memory read, memory control logic 31 may generate one or more requests that include the requested address within system memory and corresponding control information such as start-read or pre-charge commands, for example.

[018] In addition to the address and control information, the request may include error detection information such as parity information, for example. In such an embodiment,

the error detection information may include one or more parity bits which are dependent upon and protect the address and control information that is transmitted from the memory controller 30 to the memory module(s). It is noted that similar to the address and control information, the error detection information may be sent to each memory module in a

5       point-to-multipoint arrangement. Error detection generation circuit 32 may be configured to generate the error detection information. It is noted that in an alternative embodiment, the error detection information may be transmitted independently of the request. It is noted that in other embodiments, the error detection information may include other types of error detection codes such as a checksum or a cyclic redundancy code (CRC), for

10     example. Further, it is noted that in yet other embodiments, the error detection information may be an error correction code such as a Hamming code, for example. In such an embodiment, error detection circuit 130 may be configured to detect and correct errors associated with received memory requests.

15     [019] In the illustrated embodiment, system memory 40 includes memory module 0 through memory module n. Depending on the system configuration, the memory modules may be grouped into a number of memory banks such that a given number of modules may be allocated to a given range of addresses. Each signal of memory bus 35 may be coupled to each of memory modules 0 through n. Control logic (not shown in

20     FIG. 2) within each memory module may control which bank responds to a given memory request. It is noted that in an alternative embodiment, the address and command signals may be duplicated and routed among the memory modules to reduce loading effects.

25     [020] Turning to FIG. 3, a block diagram of one embodiment of a memory module is shown. Memory module 100 includes a control logic unit 110 which is coupled to sixteen memory chips, labeled MC 0-15. Memory chips 0-15 are logically divided into four banks, labeled 0-3. Memory bus 35 conveys address and control information and

data to memory module 100. The address and control signals are routed to control logic unit 110. The data path is routed directly to memory chips 0-15. Control logic unit 110 includes a buffer 120. Buffer 120 includes an error detection circuit 130. It is noted that although sixteen memory chips are shown, it is contemplated that other embodiments 5 may include more or less memory chips. Although, it is noted that four banks are described, other embodiments are contemplated in which other numbers of memory banks are used including accessing memory chips 0-15 as one bank.

[021] As described above, in one embodiment, the memory chips may be implemented 10 in DRAM. To access a location in a DRAM, an address must first be applied to the address inputs. This address is then decoded, and data from the given address is accessed. The rows and columns may be addressed separately using row address strobe (RAS) and column address strobe (CAS) control signals. By using RAS and CAS signals, row and column addresses may be time-multiplexed on common signal lines, 15 contact pads, and pins of the address bus. To address a particular memory location in a DRAM as described above, a RAS signal is asserted on the RAS input of the DRAM, and a row address is forwarded to row decode logic on a memory chip. The contents of all locations in the addressed row will then be sent to a column decoder, which is typically a combination multiplexer/demultiplexer. After row addressing is complete, a CAS signal 20 is asserted, and a column address is sent to the column decoder. The multiplexer in the column decoder will then select the corresponding column from the addressed row, and the data from that specific row/column address is placed on the data bus for use by the system.

25 [022] Control logic unit 110 receives memory requests via memory bus 35. As described above, a memory request may include address information such as the row address and the column address designated ADX, control information such as the RAS and CAS and error detection information. Each received request may be temporarily

stored in buffer 120. Control logic unit 110 may generate appropriate control signals for accessing the appropriate bank of memory chips. In the illustrated embodiment, for example, write enables (WE0, WE1, WE2, WE3), row address strobes (RAS0, RAS1, RAS2, RAS3) and column address strobes (CAS0, CAS1, CAS2, CAS3) may be

5 generated by control logic unit 110 dependent upon the received address and control information. It is noted that dependent upon the type of memory chips used (e.g. SDRAM), the control information received via memory bus 35 and generated by control logic unit 110 may include other signals (not shown).

10 [023] In addition, error detection circuit 130 generates new error detection information dependent upon the address and command information received with each request. The new error detection information is compared with the received error detection information to determine if there is an error present in the request. If an error is detected, error detection circuit 130 may transmit an error indication to memory controller 30 of FIG. 2.

15 However it is noted that in other embodiments, error detection circuit 130 may transmit the error indication to processor 20 or to a diagnostic subsystem (not shown) to indicate the presence of an error. It is noted that error detection circuit 130 may be implemented in any of a variety of circuits such as combinatorial logic, for example. It is noted that in one embodiment, the error indication may be sent from each memory module to memory controller 30 in a point-to-point arrangement, thus allowing memory controller 30 to determine which memory module has detected an error.

20  
25 [024] Depending on the configuration of system memory 40, the error may be isolated to a particular memory module, signal trace or wire. In one embodiment, the diagnostic processing subsystem may determine the cause of the error. The diagnostic processing subsystem may further isolate and shut down the failing component, or the diagnostic processing subsystem may reroute future memory requests. In other embodiments, the

diagnostic subsystem may determine the cause of the error and run a service routine which may notify repair personnel.

[025] If the current memory request is a read, error detection circuit 130 may send the  
5 error indication to memory controller 30 and control logic 110 may only send the error indication and not return any data. In response to receiving the error indication, memory control logic 31 may return a predetermined data value to processor 20 in response to receiving the error indication. Thus, in one embodiment, processor 20 may systematically abort any process which depends on that particular data. In one embodiment, the  
10 predetermined data value may be a particular data pattern that processor 20 may recognize as possibly erroneous data. In an alternative embodiment, the data may be accompanied by a bit which identifies to processor 20 that the data has an error.

[026] If the current memory request is a write, error detection circuit 130 may send the  
15 error indication to memory controller 30, thus notifying memory controller 30 that the data written to memory may have an error. In an alternative embodiment, in addition to sending the error indication to memory controller 30, error detection circuit 130 may also cause control logic unit 110 to inhibit generation of any write enable signals thus preventing data from being written into memory chips 0-15.  
20

[027] Referring collectively to FIG. 2 and FIG. 3, memory control logic 31 receives the error indication from system memory 40. In response to receiving the error indication, memory control logic 31 may store status information such as the address being written to or read from and the error indication, for example. The status information may be used in determining the cause of the error. In addition, memory control logic 31 may issue an interrupt to the diagnostic processing subsystem (not shown) or alternatively to processor 20.

[028] It is noted that in an alternative embodiment, memory control logic 31 may include a history buffer (not shown) which stores a predetermined number of past memory transactions. Thus, if error detection circuit 130 detects an error in a received request the first time that request is received, control logic 110 may inhibit writing any data to memory chips 0-15. Further, control logic 110 may send the error indication to memory control logic 31 a predetermined number of cycles after the error was detected. In response to receiving the error indication, memory control logic 31 may know how many cycles ago the error occurred. Memory control logic 31 may access the history buffer and send the correct number of past memory transactions to system memory 40. If an error is detected while resending the transactions in the history buffer, control logic 110 may inhibit generation of any write enable signals to memory banks 0-3, thus preventing data from being written into memory chips 0-15. Control logic 110 may then send the error indication to memory control logic 31 a second time. Memory control logic 31 may then send an interrupt as described previously above.

15

[029] It is noted that in one embodiment, memory bus 35 may convey address and control information in packets. In such an embodiment, the error detection information may protect the address and control information conveyed in each packet.

20

[030] However in an alternative embodiment, it is contemplated that memory bus 35 may convey address, control and error detection information in a conventional shared bus implementation. In such an embodiment, the error detection information may protect the address and control information during each address and /or clock cycle.

25

[031] Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.